

WHAT IS CLAIMED IS:

1. A network switch comprising:

two or more ports; and

a machine-readable medium embodying information indicative of instructions that when performed by the network switch result in operations comprising:

storing a value representing a determined distance to a network core for the network switch, wherein the network core represents a high bandwidth portion of a network and comprises multiple core switches having a distance to network core of zero;

configuring one of the two or more ports as a root port, in a forwarding state, that represents a least cost path to a root switch and another of the two or more ports in a blocking state, the root switch having been selected for a spanning tree for the network to center the spanning tree at the network core, the spanning tree being formed in part by said configuring; and

recovering the spanning tree, in response to a communication failure, based on periodically propagated network core distance information stored in the network switch.

2. The network switch of claim 1, wherein the periodically propagated network core distance information comprises information indicating a distance to the root switch from the network switch.

3. The network switch of claim 1, wherein recovering the spanning tree comprises:

determining whether the blocked port is operationally connected to a second switch and whether the second switch is closer to the root switch than the network switch by checking the periodically propagated network core distance information stored in the network switch;

changing the blocked port to a forwarding state if the blocked port is operationally connected to a second switch that is closer to the root switch than the network switch; and

initiating regeneration of the spanning tree if the blocked port is not operationally connected to a second switch that is closer to the root switch than the network switch.

4. The network switch of claim 1, wherein the operations further comprise, periodically:

receiving one or more frames that each identify an upstream broadcasting switch and a first distance to network core value;

storing information identifying one or more switch ports that are operational and one or more directly-connected upstream switches; and

broadcasting one or more frames that each identify the network switch and a second distance to network core value.

5. A system comprising:

a plurality of switches including core switches that reside in a network core representing a high bandwidth portion of a network, each of the core switches having a distance to network core of zero; and

machine-readable media embodying information indicative of instructions that when performed by the plurality of switches result in operations comprising:

generating a spanning tree for the network by selecting one of the core switches to be a root switch of the spanning tree; and

recovering the spanning tree in response to a communication failure based on periodically propagated network core distance information stored in at least a first switch of the plurality of switches.

6. The system of claim 5, wherein the periodically propagated network core distance information comprises information indicating a distance to the root switch.

7. The system of claim 5, wherein recovering the spanning tree comprises:
determining whether a blocked port at the first switch is operationally connected to a second switch of the plurality of switches and whether the second switch is closer to the root switch than the first switch by checking the periodically propagated network core distance information stored in the first switch;

changing the blocked port to a forwarding state if the blocked port is operationally connected to a second switch that is closer to the root switch than the first switch; and

initiating regeneration of the spanning tree if the blocked port is not operationally connected to a second switch that is closer to the root switch than the first switch.

8. The system of claim 5, wherein the operations further comprise:
periodically propagating the network core distance information downstream by broadcasting frames periodically, each frame identifying a broadcasting switch and a distance to network core value; and

updating downstream switches with the network core distance information by storing at each of the downstream switches information identifying one or more switch ports that are operational and one or more directly-connected upstream switches.

9. A method comprising:
storing a value representing a determined distance to a network core for a network switch, wherein the network core represents a high bandwidth portion of a network and comprises multiple core switches having a distance to network core of zero;

configuring a first port of the network switch as a root port, in a forwarding state, that represents a least cost path to a root switch and a second port of the network switch in a blocking state, the root switch having been selected for a spanning tree for the network to center the spanning tree at the network core, the spanning tree being formed in part by said configuring; and

recovering the spanning tree, in response to a communication failure, based on periodically propagated network core distance information stored in the network switch.

10. The method of claim 9, wherein the periodically propagated network core distance information comprises information indicating a distance to the root switch from the network switch.

11. The method of claim 9, wherein recovering the spanning tree comprises:
determining whether the blocked port is operationally connected to a second switch and whether the second switch is closer to the root switch than the network switch by checking the periodically propagated network core distance information stored in the network switch;

changing the blocked port to a forwarding state if the blocked port is operationally connected to a second switch that is closer to the root switch than the network switch; and

initiating regeneration of the spanning tree if the blocked port is not operationally connected to a second switch that is closer to the root switch than the network switch.

12. The method of claim 9, further comprising, periodically:
receiving one or more frames that each identify an upstream broadcasting switch and a first distance to network core value;

storing information identifying one or more switch ports that are operational and one or more directly-connected upstream switches; and

broadcasting one or more frames that each identify the network switch and a second distance to network core value.

13. An article comprising a machine-readable medium embodying information indicative of instructions that when performed by one or more machines result in operations comprising:

storing a value representing a determined distance to a network core for a network switch, wherein the network core represents a high bandwidth portion of a network and comprises multiple core switches having a distance to network core of zero;

configuring a first port of the network switch as a root port in a forwarding state that represents a least cost path to a root switch and a second port of the network switch in a blocking state, the root switch having been selected for a spanning tree for the network to center the spanning tree at the network core, the spanning tree being formed in part by said configuring; and

recovering the spanning tree, in response to a communication failure, based on periodically propagated network core distance information stored in the network switch.

14. The article of claim 13, wherein the periodically propagated network core distance information comprises information indicating a distance to the root switch from the network switch.

15. The article of claim 13, wherein recovering the spanning tree comprises:
determining whether the blocked port is operationally connected to a second switch and whether the second switch is closer to the root switch than the network switch

by checking the periodically propagated network core distance information stored in the network switch;

changing the blocked port to a forwarding state if the blocked port is operationally connected to a second switch that is closer to the root switch than the network switch; and

initiating regeneration of the spanning tree if the blocked port is not operationally connected to a second switch that is closer to the root switch than the network switch.

16. The article of claim 13, wherein the operations further comprise, periodically: receiving one or more frames that each identify an upstream broadcasting switch and a first distance to network core value;

storing information identifying one or more switch ports that are operational and one or more directly-connected upstream switches; and

broadcasting one or more frames that each identify the network switch and a second distance to network core value.